Technical Note: Locality Sensitive Reflections (LSR)

Background (key facts)

- **Locality Sensitive Hashing (LSH)**
  - For angular similarity:
  - $s(x)=\text{sign}(Rx)$, $R \in \text{Real}^{k \times n}$
  - Partitions space by random hyperplanes; nearby vectors share sign patterns with high probability.
- **Reflections as involutions**
  - A diagonal matrix with ±1 entries:

    $D = \text{diag}(\pm 1), \quad D^2 = I$

  - Represents a reflection across coordinate hyperplanes.
- **Structured transforms (Hadamard-type)**
  - Products like HD (with orthogonal H) act as rotations/mixers. Two reflections equals a rotation. $H^2=I$ (second reflection)
  - Fast Johnson–Lindenstrauss transforms exploit this for approximate isometries.
- **Piecewise-linear / piecewise-isometric maps**
  - Space partitioned into regions; each region applies a fixed linear or orthogonal map.
  - Common in ReLU networks (piecewise linear) and dynamical systems (piecewise isometries).

Definition

A **Locality Sensitive Reflection (LSR)** is a data-dependent involution:

$D(x)=\text{diag}(\text{sign}(Rx))$

Induced operator:

$T(x) = D(x)x$ or $T(x) = WD(x)x$

- D(x) is constant within each LSH cell.
- Globally, T is **piecewise orthogonal** (or piecewise linear if followed by W).

Properties

- **Involution (per region)**:
  $D(x)^2 = I$
- **Local linearity**:
  On each region Region_i, $D(x) = D_i \Rightarrow T(x)$ is linear.
- **Global nonlinearity**:
  Discontinuities occur on hyperplane boundaries.
- **Geometry-aware partitioning**:
  Regions defined by random hyperplanes preserve angular locality.
- **Cheap evaluation**:
  - Hash: O(kn)
  - Apply: O(n)

Composition with orthogonal transforms

$x \mapsto HD(x)x$

- Each region induces a fixed orthogonal map HD_i.
- Globally: **piecewise rotation/reflection field**.

Iterative form:

$x_{t+1} = HD(x_t)x_t$

- Defines a **discontinuous dynamical system** on the sphere.

Relation to existing constructs

- **Random projections / JL transforms**: fixed D → randomized rotation
- **Hypernetworks / dynamic layers**: W(x) depends on input
- **Mixture-of-experts**: region-dependent linear maps
- **Binary neural networks**: sign-based nonlinearity

- **Householder-based methods**: compositions of reflections generate orthogonal maps

LSR can be viewed as:

> a **quantized, hash-driven, data-dependent orthogonal operator**

Limitations

- **Non-differentiability** (sign function)
- **Low information content** (1-bit per projection)
- **No magnitude gating** (only sign flips)
- **Region aliasing** (distinct inputs share same reflection)

Variants

- **Smoothed LSR**:

  $D(x)=\mathrm{diag}(\tanh(\alpha Rx))$

- **Multi-hash LSR**:

  $$\sum_i W_i D_i(x) x$$

- **Decoupled form**:

  $W(D(x)Vx)$

Further study keywords

- *piecewise isometries*
- *fast Johnson–Lindenstrauss transform*
- *Householder transformations*
- *orthogonal neural networks*
- *binary neural networks / sign activations*
- *mixture-of-experts / conditional computation*
- *hypernetworks / dynamic linear layers*
- *normalizing flows (orthogonal / coupling layers)*

- *angular locality sensitive hashing*
- *random feature maps*
- *geometric algebra (reflections and rotors)*

Summary

Locality Sensitive Reflections convert LSH from a passive indexing tool into an **active, data-dependent operator**, yielding:

- piecewise orthogonal transformations
- geometry-aware nonlinear mappings
- extremely low-cost adaptive behavior

They occupy a middle ground between:

- fixed random projections
- and fully learned nonlinear networks.